

Myopic Agency

JOHN Y. ZHU¹

April 1, 2013

Abstract

I consider a repeated principal-agent setting in which the agent repeatedly chooses between hidden “long-term” and “short-term” actions. Relative to the long-term action, the short-term action boosts output today but hurts output tomorrow. The optimal contract inducing long-term actions provides a fresh perspective on upward sloping pay scales, severance pay and high-watermark contracts. The myopic agency model is unlike traditional agency models that ignore action persistence. I show ignoring action persistence can significantly distort optimality. For example, traditional optimal contracts typically reward something like high average output whereas the myopic agency optimal contract rewards sustained high output.

JEL Codes: C73, D86, J33, J41, M52

¹zhuyiran@wharton.upenn.edu, The Wharton School, University of Pennsylvania. I thank Daniel Neuhann for excellent research assistance and Richard Kihlstrom for providing valuable guidance throughout the development of the paper. I am grateful to Robert Anderson, Alex Edmans, Mike Fishman, Itay Goldstein, Ben Hermalin, Steven Matthews, Andy Postlewaite, Andy Skrzypacz, Rob Stambaugh and Bilge Yilmaz for comments and advice. Lastly I would like to thank the seminar participants at the University of Pennsylvania, 2012 NBER/NSF/CEME Conference on Mathematical Economics and General Equilibrium Theory, and Northwestern Finance for helpful feedback.

1 Introduction

How can a firm’s owners prevent a manager from taking hidden actions that look good today but hurt long-run profitability? The large literature on moral hazard has surprisingly little to say about this problem. I call this agency problem faced by the owners “myopic” agency. In this paper, I investigate myopic agency in a dynamic principal-agent setting. At each date, the agent takes a hidden action that has persistent effects on firm performance. There are two actions: long-term and short-term. The long-term action maintains a certain benchmark level of expected output. The short-term action causes current expected output to rise above the benchmark and future expected output to drop below. The drop is assumed to be sufficiently large relative to the rise so that the principal prefers the long-term action. I then explicitly characterize and study the optimal incentive contract that always induces the long-term action from the agent.

To get a feel for the potential pitfalls of contracting under repeated myopic agency, consider the problem of trying to induce a manager to take the long-term action today. The usual view of moral hazard, familiar in an insurance setting, tells us to reward good outcomes and punish bad outcomes. But rewarding high output today will only encourage the manager to take the short-term action. A more sensible strategy is to wait until tomorrow (when the long-term effects of today’s action have been realized) and reward the agent only if high output is produced then. While this strategy works in a one-shot model of myopic agency, in a dynamic setting, this arrangement only serves to pass today’s agency problem onto tomorrow. Facing such a contract with delayed rewards and punishments, a sophisticated manager will simply behave today and wait until tomorrow to take the short-term action.

In what settings will an agent repeatedly face a long-term/short-term decision? A natural one is a firm’s R&D department. Channeling resources to R&D helps increase the long-term profitability of the firm. The short-term action of not fostering R&D will help boost profits and dividends today, but may cause the firm to become obsolete in the future. More generally, settings where the manager must make investment decisions are vulnerable to myopic agency if the quality of the decisions cannot be well-monitored or understood.

For decades, Kodak employed a business model that involved selling inexpensive cameras and then generating a large profit margin through its lucrative film business. Naturally, most of the investment went into film as opposed to cameras. In the 90s as the digital revolution was in full swing, Kodak faced a difficult decision between doubling down on film or switching its focus to cameras. In Kodak’s case, the agency problem wasn’t that the firm underinvested in R&D. The problem was that Kodak chose to stick with film. Throughout the 90s, the company suffered very little from this decision. Even in the late 90s, the firm sold about 1 billion rolls of film each year. However, by 2011, sales had sunk to 20 million.

Financial markets are also rife with myopic agency problems. Subprime lending is a perfectly legitimate long-term action, but doing it prudently requires effort to carefully vet the borrowers and assess the complex associated risks. Without the proper incentives, an agent may engage in indiscriminate lending under terms overly favorable to the borrower. While deviating to this short-term action is an easy way to inflate business today, the long-term effects can be disastrous. Hedge funds are also susceptible to myopic agency. Here, the desired long-term action involves the manager’s exerting effort to turn his innate skill into

generating alpha. The undesirable short-term action can be employing strategies that are essentially equivalent to writing a bunch of puts, exposing the fund to significant future tail risk.

A common property in these examples is that the moral hazard problem is played out over time. It cannot be properly modeled when actions in date t affect the outcome only in date t . When actions are non-persistent, rewarding high output, as a general rule, helps alleviate the agency problem. However, in a myopic agency setting, rewarding high output can exacerbate the agency problem just as much as alleviate it. Divorced from the past and the future, a stochastic output today communicates very little information about the agent's decisions. As a result, the principal must pay careful attention to the pattern of production across time. The structure of the optimal dynamic contract will reflect this requirement.

The first salient feature of the optimal contract is how it decides when the agent is doing well. A typical optimal dynamic contract will formulate an endogenous measure(s) of good performance based on the history of outcomes. It will then use this measure to determine the spot contract that the agent faces today. In the myopic agency model, high output is a priori an ambiguous signal of performance. High output today is a positive signal only when tomorrow's output is high. Otherwise, it really looks like the agent took the short-term action today. Similarly, tomorrow's high output is a positive signal only if the day after tomorrow's output is high. Thus, the surest indicator that the agent is behaving is if there is an *unbroken string of high outputs*. This means the optimal contract cares foremost about consistency. Specifically, the endogenous good performance measure of the optimal contract tracks the number of consecutive high output dates leading up to today.

Consistency is widely extolled by employers. Orthodox thinking is that consistency is valuable because it signals ability or because it provides a stable setting allowing investment opportunities and worker morale to flourish. The optimal contract shows that there is an *agency-driven* contribution to the value of consistency.

On the flip side, because the optimal contract values consistency, low output after a long history of high output can have a seemingly disproportionate effect on the agent's standing with the principal. This creates a cliff-like arrangement where, the longer the high output streak, the higher the agent ascends on the contractual ladder but also the farther he falls when the streak is broken. I show that this cliff arrangement is a natural way of dealing with the double-edged sword of rewarding high output in a myopic agency setting.

The cliff arrangement helps explain why upward sloping pay scales feature prominently in many employment contracts.² It provides some context for why high-watermark contracts are popular in hedge funds, given the particular characteristics of their returns. And it predicts that incentive level and pay-to-performance sensitivity should both increase after good performance.³

When termination for poor performance is introduced, the optimal contract also features severance pay - even large, lump sum severance pay. The intuition comes straight from the myopic agency problem and is not sensitive to the particular details of the optimal contract.

²See Lazear (1981) for an alternative argument for upward sloping pay scales involving the backloading intuition and firm-side incentive problems.

³There does not seem to be an empirical consensus on this conjecture. In general, the potential connections between performance and pay-to-performance sensitivity seems to have received little empirical attention.

For a simple explanation, consider the following setup: Suppose the agent is at a point in his employment where a low output today will mean termination. If termination is very costly, or equivalently if staying employed is very lucrative, then it is imperative to produce high output today and the short-term action becomes quite attractive. One possible way to discourage short-termism today despite the strong temptation is to promise the agent a very large reward for high output tomorrow. The agent, understanding that he has little chance of collecting this reward tomorrow if he takes the short-term action today, will choose the long-term action instead.

But a large future reward is not the only way to fight short-termism in the face of an imminent termination threat. The alternative approach, arguably more natural, is to simply make termination less painful for the agent. Severance pay does precisely this. Moreover, the larger the severance pay and the more lump sum it is, the less painful termination becomes.

The two separate approaches to fighting short-termism mean the principal faces a trade-off: he can either pay the agent more today following termination or pay the agent more tomorrow following high output. The tradeoff favors paying the agent today and is what makes severance pay a robust feature. This result is not by accident. Paying the agent more tomorrow for high output makes taking the short-term action tomorrow more tempting, which means that the principal will also have to pay the agent more the day after tomorrow for high output, and the day after the day after tomorrow, and so on. Thus by choosing severance pay the principal avoids an infinite cascading escalation of pay in the future.

1.1 A Comparative Discussion

The optimal dynamic contract under myopic agency differs in significant ways from many of those that arise under traditional, non-persistent moral hazard. Unlike the consistency-based measure of good performance used by the myopic agency optimal contract, traditional dynamic contracts typically have a measure that resembles counting total high output dates net low output dates. The aggregation is insensitive to the specific timing of performance. This type of approximately order-independent aggregation of performance over time has been the dominant measure of performance quality in the dynamic contracting literature since the subject's inception. Early examples include the cumulated performance S of Radner (1981, 1985) and the linear aggregator of Holmstrom and Milgrom (1987).

More generally, myopic agency forces us to change how we think about what makes a contract good or bad. Traditional contract theory's basic tenet is: if you want to get the agent to do the right thing, you have to provide him with enough incentives. This value reflects the traditional IC-constraint, which is a lower bound on incentives. The implication being, a bad contract is one whose incentives are too small. The myopic agency IC-constraint is, instead, concerned with balancing incentives across time. A bad contract isn't necessarily one that has small incentives, but rather one that has, in some sense, inappropriate incentives.

This paradigm shift in the intuition for incentive-compatibility also has theoretical implications. In most traditional dynamic arrangements, the state variable is the agent's promised value. Seminal papers employing the recursive technique over promised values include Green (1987), Spear and Srivastava (1987) and Abreu, Pearce and Stacchetti (1990). Under myopic agency, the optimal contract's state variable is the incentive level: the *sensitivity* of the

agent's promised value to output. Therefore, the optimal contract of this paper exemplifies a new species of arrangements where promised value plays a genuinely secondary role.

In the paper, I formally compare the myopic agency optimal contract and the contract the principal would use if he decided to abstract away from the persistent nature of the actions. I call the latter contract the traditional contract. It is the optimal contract of the appropriate non-persistent version of the model. I find that the traditional contract isn't terrible in the sense that if it is used in the actual myopic agency setting in lieu of the optimal contract, the agent will still take the long-term action. However, it always over-incentivizes the agent relative to the optimal contract and is therefore inefficient.

The traditional contract has a stationary incentive level. In contrast, the optimal contract's expected incentive level increases in a concave, convergent way over time. This type of behavior is also documented in Gibbons and Murphy (1992), which provides an alternative justification that rests on career concerns and requires risk aversion. The incentive escalation result of this paper is compatible with - but not dependent on - risk aversion. Another difference is that the incentive escalation in my optimal contract exhibits history-dependent dynamics, escalating only after high output.

While this paper explores how to induce the long-term action through contracts, there is a related literature that focuses on why managers oftentimes take a variety of short-term actions in equilibrium. See, for example, Stein (1988, 1989). Another related literature deals with innovation. The process generates a dynamic not unlike the one produced by the long-term action. Manso (2012) embeds such a two-date innovation problem within a principal-agent framework. Edmans, Gabaix, Sadzik and Sannikov (2012) consider a model of dynamic manipulation that allows the agent to trade off, on a state-by-state basis, future and present performance. Varas (2013) considers a model of project creation where the principal faces a compensation problem similar to the one in this paper. By rewarding the agent for the timely completion of a good project whose quality is hard to verify, the principal might inadvertently induce the agent to cheat and quickly produce a bad project.

This paper is most closely related to Holmstrom and Milgrom (1991). Recall, they observe that if the agent has two tasks A and B , the incentives of A may exert a negative externality on that of B . In my model, one can think of the task of managing the firm today as task A and managing the firm tomorrow as task B . And just as in Holmstrom and Milgrom (1991), if incentives today are too strong relative to those of tomorrow, the agent will take the short-term action, which favors the firm today and neglects the firm tomorrow. Now, Holmstrom and Milgrom use this to explain why contracts often have much lower-powered incentives than what the standard single-task theory might predict. In my paper things are further complicated by the dynamic nature of the model. Specifically, today's task B will become tomorrow's task A . Each date's task is both task A and task B depending on the frame of reference. Therefore, the conclusion in my model is not that incentives should be low-powered, but that incentives start low and optimally escalate over time.

My paper is also part of a small literature on persistent moral hazard. An early treatment by Fernandes and Phelan (2000) provides a recursive approach to computing optimal contracts in repeated moral hazard models with effort persistence. Jarque (2010) considers a class of repeated persistent moral hazard problems that admit a particularly nice recursive formulation: those with actions that have exponential lagged effects. She shows that under




Figure 1: Transition function between the good and bad states.

a change of variables, models in this class translate into traditional non-persistent repeated moral hazard models. Her work can be interpreted as a justification for the widely used modeling choice of ignoring effort persistence in dynamic agency models.

My paper considers, in some sense, the opposite type of persistence to that of Jarque (2010). Here, the ranking of the actions is flipped over time. Today: short-term $>$ long-term; future: long-term $>$ short-term. With this type of persistence, results become noticeably different from those of the non-persistent class.

The rest of the paper is organized as follows: Section 2 introduces the basic repeated myopic agency model. I recursively characterize and solve for the optimal contract. Section 3 interprets the optimal contract. The novel performance measure and the cliff arrangement emerge. Comparisons with high watermark contracts and non-persistent optimal contracts are made. Section 4 deals with incentive escalation and the option to terminate. Upward sloping pay scales and severance pay emerge as optimal arrangements. Section 5 concludes.

2 Repeated Myopic Agency

A principal contracts an agent to manage a firm at dates $t = 0, 1, 2, \dots$. At each date t , the firm can be in one of two states: $\sigma_t = good$ or bad . If $\sigma_t = good$, then the agent can apply one of two hidden actions: a long-term action $a_t = l$ or a short-term action $a_t = s$. If the agent applies the long-term action, then the firm remains in the good state: $\sigma_{t+1} = good$. If the agent applies the short-term action, then $\sigma_{t+1} = good$ with probability $\pi < 1$, and $\sigma_{t+1} = bad$ with probability $1 - \pi$. If $\sigma_t = bad$, then there is no action choice and the state reverts back to $good$ at the next date. See Figure 1.

Actions and states are hidden from the principal, who can only observe output. At each date t , the firm produces either high output $X_t = 1$ or low output $X_t = 0$. If $\sigma_t = good$ and $a_t = l$ then the probability that the firm produces high output at date t is $p < 1$. If $\sigma_t = good$ and $a_t = s$ then the firm produces high output for sure at date t . If $\sigma_t = bad$ then the firm produces low output for sure at date t . I assume that $\sigma_0 = good$.

Notice, if the agent always takes the long-term action, then the firm is always in the good state and there is always a probability p of high output. A deviation today to the short-term action boosts expected output today by $1 - p$ and lowers expected output tomorrow by $Q := (1 - \pi)p$. I assume that $1 - p < \beta Q$ where $\beta \in (0, 1)$ is the intertemporal discount factor. This assumption says that the gain today from taking the short-term action is outweighed by the present discounted loss tomorrow.

Definition of a Contract.

At each date t , the principal may make a monetary transfer $w_t \geq 0$ to the agent. Note, each w_t can depend on the history of outputs up through date t . However, w_t cannot depend on the unobservable action nor the state. At each date t , the principal may also recommend an action a_t to be taken provided $\sigma_t = \text{good}$. A contract is a complete transfer and action plan $w = \{w_t\}, a = \{a_t\}$. The principal's utility is $\mathbf{E}_a \left[\sum_{t=0}^{\infty} \beta^t (X_t - w_t) \right]$ and the agent's utility is $\mathbf{E}_a \left[\sum_{t=0}^{\infty} \beta^t (w_t + b1_{a_t=s}) \right]$. I assume selecting the short-term action provides private benefit $b > 0$ to the agent. I also assume that the agent has an outside option worth 0. This assumption eliminates the participation constraint.

Assumption (A). *The principal always wants to induce the agent to take the long-term action.*

The action sequence taken by the agent should, in principle, be determined as part of the optimal contracting problem. However, I show that requiring the agent to take the long-term action is without loss of generality under certain parameterizations of the model. See Corollary to Theorem 1. Moreover, the solution to the sustained long-term action case will serve as an important benchmark for future analyses of the unconstrained optimal contracting problem. Given Assumption (A), I define the optimal contracting problem to be the following constrained maximization:

$$\begin{aligned} & \max_{\{w_t \geq 0\}_{t=0}^{\infty}} \mathbf{E}_{\{a_t=l\}_{t=0}^{\infty}} \left[\sum_{t=0}^{\infty} \beta^t (X_t - w_t) \right] \\ \text{s.t.} \quad & \{a_t = l\}_{t=0}^{\infty} \in \arg \max_{\{a_t\}} \mathbf{E}_{\{a_t\}} \left[\sum_{t=0}^{\infty} \beta^t (w_t + b1_{a_t=s}) \right] \end{aligned}$$

Throughout the paper, if there are multiple optimal arrangements, I focus on the one that pays the agent earlier. This ensures that the optimal contract is approximately robust to small perturbations to the agent's discount factor that make him more impatient than the principal.

Let \mathcal{H}_t denote the set of all binary sequences of length $t + 1$. \mathcal{H}_t is the set of histories of firm outputs up through date t . In general, the agent's promised value depends on the history of outputs up through yesterday as well as today's state. So for each $h_{t-1} \in \mathcal{H}_{t-1}$ and state σ_t , define $W_t(h_{t-1}, \sigma_t)$ to be the agent's date t promised value given all the relevant information:

$$W_t(h_{t-1}, \sigma_t) = \mathbf{E}_a \left[\sum_{i=t}^{\infty} \beta^{i-t} (w_i + b1_{a_i=s}) \mid h_{t-1}, \sigma_t \right]$$

In general, the promised value is unknown to the principal since states are hidden. However, the paper restricts attention to only those contracts where the agent takes the long-term action all the time and the state is always good. Therefore, on the equilibrium path, it is well-defined to speak of $W_t(h_{t-1})$, which only depends on the publicly observable h_{t-1} and is known to the principal.

In the traditional approach to solving dynamic contracting problems, the agent's promised value is the key state variable and the entire recursive formulation of optimality is built around it. This will not be the case in the myopic agency model. For each history h_{t-1} , define $\Delta_t(h_{t-1}) := [w_t(h_{t-1}1) + \delta W_{t+1}(h_{t-1}1)] - [w_t(h_{t-1}0) + \delta W_{t+1}(h_{t-1}0)]$. As a function over \mathcal{H}_{t-1} , Δ_t is a random variable representing the date t incentive level of the contract. It turns out that Δ_t is the natural state variable of the myopic agency model. Thus, the optimal contract will be recursive over Δ_t , not W_t . The centrality of Δ_t is due to the incentive-compatibility condition:

Lemma 1. *A contract always induces the long-term action if and only if at each date t and after each history $h_{t-1} \in \mathcal{H}_{t-1}$*

$$\Delta_{t+1}(h_{t-1}1) \geq \varepsilon(\Delta_t(h_{t-1})) := \frac{(1-p)\Delta_t(h_{t-1}) + b}{\beta Q} \quad (1)$$

Proof. See Appendix. □

Even though incentive-compatibility involves comparing on- and off-equilibrium continuation values, notice that Equation (1) only involves on-equilibrium continuation values. In the Appendix, I explain how this is possible by first showing that the agent's off-equilibrium continuation value can be expressed as a function of on-equilibrium continuation values following a one-shot deviation. Equation (1) is precisely the condition that prevents such one-shot deviations. I then show that checking for one-shot deviations is sufficient.

The IC-constraint is a lower bound for the incentives tomorrow as a function of the incentives today. Or equivalently, it is an upper bound on incentives today as a function of incentives tomorrow. Either way, the absolute levels of incentives do not matter so much, what matter are the relative levels of incentives over time. The greater the incentives are today, the more tempting it is to take the short-term action today. Therefore, the incentives tomorrow must also keep pace to ensure that the agent properly internalizes the future downside of taking the short-term action today. Similarly, if the private benefit b is large, then it is again tempting to take the short-term action today. So the lower bound on tomorrow's incentives is also an increasing function of b .

Notice the IC-constraint cares about tomorrow's incentives following high output today but ignores tomorrow's incentives following low output today. This is because if the agent actually deviates and takes the short-term action, low output today never occurs. Since the incentive level after low output is immaterial under a short-term action deviation, it is not included in the IC-constraint.⁴

⁴What would happen if, more generally, taking the short-term action today only increased the probability of high output today to $p + q \leq 1$? The IC-constraint would take the more general form: $(p + q)\Delta_{t+1}(h_{t-1}1) + (1 - p - q)\Delta_{t+1}(h_{t-1}0) \geq \frac{q\Delta_t(h_{t-1}) + b}{\beta Q}$. Clearly, the generalized IC-constraint does care about the incentives tomorrow following low output today. However, the weighting of tomorrow's incentives is under the *counterfactual* measure generated by taking the short-term action today. Relative to the actual measure, the counterfactual measure underweights low output today. As a result, the generalized IC-constraint still underemphasizes the incentives tomorrow following low output today. Numerical simulations show that the optimal contract under the generalized IC-constraint is qualitatively similar to the optimal contract under Equation (1).

Assumption (B). *The agent can freely dispose of output before the principal observes the net output.*

Now, if a contract's promised value is decreasing in output, the agent will simply engage in free disposal after high output and mimic low output. Thus, Assumption (B) amounts to a monotonicity of promised value requirement and implies $\Delta_t \geq 0$ for all t .

Definition. *For any $\Delta \geq 0$, the term Δ -contract will mean a contract whose initial incentive level is Δ . Define $V(\Delta)$ to be the cost to the principal of the optimal Δ -contract.*

Since I will show that Δ is the state variable, the function $V(\Delta)$ will serve as the "optimal value function." The optimal contracting problem then boils down to finding and solving the appropriate Bellman equation characterizing $V(\Delta)$.

Lemma 2. *$V(\Delta)$ is convex and the optimal contract is randomization-proof. $V(\Delta)$ is also weakly increasing, so the optimal 0-contract is the optimal contract.*

Proof. See Appendix. □

I now give an intuitive explanation of why the optimal contract can be recursively characterized over Δ . This provides a non-formal derivation of the Bellman equation characterizing $V(\Delta)$.

Fix an arbitrary Δ and consider the optimal Δ -contract. Since the date 0 incentive level of this contract is Δ , the IC-constraint implies that date 1's incentive level following high output at date 0 must be at least $\varepsilon(\Delta)$. Therefore, viewed as a contract in its own right, this date 1 continuation contract is some Δ' -contract where $\Delta' \geq \varepsilon(\Delta)$. Since the contract under consideration is the least costly one among all Δ -contracts, it stands to reason that its date 1 continuation contracts must also be least costly subject to incentive constraints. Thus, the Δ' -contract following high output must be the least costly contract with initial incentive level $\geq \varepsilon(\Delta)$. Since V is weakly increasing, without loss of generality, the contract must be the optimal $\varepsilon(\Delta)$ -contract with cost $V(\varepsilon(\Delta))$. Using the same logic, since there are no incentive constraints following low output, the date 1 continuation contract following low output is the optimal 0-contract with cost $V(0)$.⁵

Let $w(1)$ and $w(0)$ denote the high and low output payments to the agent at date 0 in the optimal Δ -contract. So far, I have shown that $V(\Delta) = p(w(1) + \beta V(\varepsilon(\Delta))) + (1 - p)(w(0) + \beta V(0))$. Once $w(1)$ and $w(0)$ are pinned down, the recursive characterization is complete.

Given the above expression for $V(\Delta)$, the date 0 incentive level is $(w(1) + \beta V(\varepsilon(\Delta))) - (w(0) + \beta V(0))$, which must be Δ by assumption. A little algebra shows that the minimal $w(1)$ and $w(0)$ satisfying the condition are:

$$\begin{aligned} w(1) &= \left(\Delta - \beta [V(\varepsilon(\Delta)) - V(0)] \right)^+ \\ w(0) &= \left(\beta [V(\varepsilon(\Delta)) - V(0)] - \Delta \right)^+ \end{aligned} \tag{2}$$

⁵For a non-generic set of parameter values, $V(\Delta)$ is flat for an interval of Δ values. When this occurs, tomorrow's continuation contract following low output today can have an initial incentive level greater than 0 without diminishing today's payments to the agent. Since $V(\Delta)$ will eventually be characterized explicitly, we will see exactly what these non-strictly increasing versions of $V(\Delta)$ look like. As a result, the non-robust versions of the optimal contract in these cases can be easily described.

I can now express $V(\Delta)$ exclusively as a function of Δ , $V(0)$ and $V(\varepsilon(\Delta))$. This is the Bellman equation characterizing $V(\Delta)$. Solving the equation formally solves the optimal contracting problem.

Theorem 1. *The Markov law of the state variable Δ for optimal Δ -contracts is*

$$\Delta \rightarrow \begin{cases} \varepsilon(\Delta) & \text{following high output} \\ 0 & \text{following low output} \end{cases}$$

The optimality function $V(\Delta)$ satisfies the following explicitly solvable Bellman equation:

$$V(\Delta) = \max\{\beta V(\varepsilon(\Delta)) - (1-p)\Delta, p\Delta + \beta V(0)\} \quad (3)$$

The solution is a piecewise linear, weakly increasing, convex function.

Proof. See Appendix. □

$V(\Delta)$ not only implies the structure of the optimal contract but more generally, the structure of the optimal Δ -contract for any $\Delta \geq 0$. Viewed as a Markov process, Δ_t has stationary measure μ over the discrete set $\{\varepsilon^N(0)\}_{N=0}^{\infty}$. The stationary measure is $\mu(\varepsilon^N(0)) = (1-p)p^N$.

Corollary. *Fixing all other parameters, for all sufficiently small private benefit b , always inducing the long-term action is optimal.*

In the proof of Theorem 1, I show that the cost of a Δ -contract is proportional to the size of the private benefit. So as a function, $V(\Delta)$ is homogenous of degree 1 in b . After any history, the cost of always inducing the long-term action from now on is bounded above by $V(\varepsilon^\infty(0)) = V(b/(\beta Q - (1-p)))$. Thus, as b tends to zero, so will the upper bound. In particular, the upper bound will be smaller than the surplus generated from taking the long-term action a single time.

In the next section I will explicitly write down a representative solution to the Bellman equation and analyze the implied optimal contract.

3 A Representative Optimal Contract

For a large set of parameter realizations, including all realizations where $Q \geq 6 - 4\sqrt{2} \approx .343$, $V(\Delta)$ is a two piece piecewise linear function:⁶

$$V(\Delta) = \begin{cases} V(0) + m(p)\Delta & \text{if } \Delta \in [0, x_1] \\ \beta V(0) + p\Delta & \text{if } \Delta \in [x_1, \infty) \end{cases} \quad (4)$$

⁶The condition on Q is fairly mild. Since I assume in the model $1-p < \beta Q$ and since by definition $Q \leq p$, it must be that $p > \frac{1}{2}$. Therefore, the condition on Q does not further restrict the domain of p . Moreover, the closer Q is to p , the more flexibility one has in setting the other parameter, β . Formally, $V(\Delta)$ is a two piece piecewise linear function if and only if $\beta \leq \frac{Q-p(1-p)}{(p-Q)(1-p)}$.

where

- $V(0) = \frac{1}{1-\beta^2} \frac{pb}{Q}$
- $m(x) = \frac{1-p}{Q}x - (1-p)$
- $x_1 = \frac{1}{1+\beta} \frac{pb}{Q-p(1-p)}$

The optimal contract delivers expected value $V(0)$ to the agent. As a function of incentives, today's spot contract is

$$w(1) = \begin{cases} 0 & \text{if } \Delta \in [0, x_1] \\ \frac{Q-p(1-p)}{Q} [\Delta - x_1] & \text{if } \Delta \geq x_1 \end{cases} \quad (5)$$

$$w(0) = \begin{cases} \frac{Q-p(1-p)}{Q} [x_1 - \Delta] & \text{if } \Delta \in [0, x_1] \\ 0 & \text{if } \Delta \geq x_1 \end{cases} \quad (6)$$

Notice if $Q = p$, which is a case covered by the representative example, $V(\Delta)$'s first piece becomes flat. This is the only case when $V(\Delta)$ is not strictly increasing everywhere. When $Q = p$, the incentive level tomorrow can take any value between 0 and x_1 following low output today without lowering today's payment to the agent. However, since any perturbation to the value of Q will cause $V(\Delta)$ to become strictly increasing, these alternate forms are not robust.

3.1 How and Why the Optimal Contract Works

Formally, the state variable of the optimal contract is Δ . In practice, it suffices to count the number N of consecutive high output dates leading up to today. The performance indicator N completely determines today's spot contract, incentives and continuation contract.

Thus, the optimal contract cares about *sustaining* high output and not simply generating a large number of high output dates.

As a function of N , today's incentive level is $\varepsilon^N(0)$ and is increasing concave. The limiting incentive level is

$$\bar{\Delta} := \varepsilon^\infty(0) = \frac{b}{\beta Q - (1-p)} \quad (7)$$

As a function N , today's spot contract is also determined and its pay-to-performance sensitivity is also increasing concave.

Consider the payment to the agent for high output when his performance indicator is N . Viewed as a function over N , this reward schedule is zero for low values of N and then becomes a concave, convergent function. Thus, when the agent's performance is poor, he is not rewarded for high output. Then there is an intermediate stage when the agent is facing small but rapidly increasing rewards. Eventually, rewards for high output level off at a large sum for sufficiently good performance.⁷

⁷In this particular class of examples, low values of N means $N = 0$. In general, the initial period of no high output reward may last longer but not less than one date.

The performance indicator and reward schedule arrangement contains the basic intuition for why the optimal contract works. When the principal sets out to write a contract for the agent, he faces a dilemma. On the one hand, he needs to reward the agent for good performance since he is asking the agent to perform a productive yet costly hidden task. On the other hand, the principal realizes that a large reward for high output, paradoxically, tempts the agent to deviate since the short-term action guarantees the agent the reward. The optimal contract’s performance indicator-reward schedule arrangement is tailored to resolve this dilemma.

Re-imagine the indicator-reward arrangement as a path up a cliff. At the bottom of the cliff, the path is barren. At the top, the path is lined with big carrots. The path represents the domain of the performance indicator and the carrots comprise the reward schedule.

Initially, the agent is at the bottom of the cliff. The performance indicator determines his progress along the path. In particular, every high output moves the agent forward one step. His goal is to consume the big carrots at the top. To reach them, the agent must take the long-term action. Scaling the cliff requires sustained progress. Taking the short-term action today may help the agent progress today but it will not help him progress over and over again.

Once the agent has reached the top of the cliff, he can begin to consume the big carrots. The agent’s hidden actions represent two approaches to consumption. He can take sensible bites by continuing to take the long-term action or he can gorge by taking the short-term action.

Will the agent gorge? Suppose he does. Then one consequence of this choice is that the agent is likely deprived of a big carrot tomorrow. For example, if $Q = p$, then gorging today ensures no consumption tomorrow. But this is a relatively minor loss compared to the other one looming on the horizon. Recall the agent’s progress is governed by a performance indicator that cares only about sustained high output. Therefore, when the likely low output event is realized tomorrow, the streak is broken and the indicator drops down to zero.

And the agent falls off the cliff.

3.2 The Cliff Arrangement and High-Watermark Contracts

As the cliff analogy demonstrates, the optimal contract motivates the agent in two mutually reinforcing ways. Initially, the backloaded nature of the reward schedule induces the agent to take the long-term action. This, in and of itself, is unremarkable. Many optimal dynamic contracts have some form of backloading of rents.

The novelty comes when the backloaded reward schedule is mapped against the optimal contract’s novel performance indicator N . Because N takes precipitous drops, this combination creates a “contractual cliff” that provides the second way to motivate the agent - when performance has already reached a high level and the previously backloaded payments have come to the fore. At this point, the fear of falling off the cliff and starting all over again serves as an effective deterrent to short-termism. This cliff arrangement differs from typical optimal arrangements seen in traditional non-persistent dynamic moral hazard settings. In those settings, optimal contracts typically care about something that roughly approximates

high output dates net low output dates. This inattention to the timing of high output effectively means there’s no cliff.

How serious is this omission in a myopic agency setting?

Remark. *Suppose the principal eliminates the optimal contract’s cliff entirely. At each date, the agent is simply given the limiting spot contract as the performance indicator goes to infinity. The limiting spot contract’s pay-to-performance sensitivity is larger than that of any of the spot contracts used in the optimal contract. Despite having more “skin in the game,” the agent now never chooses to take the long-term action.*

To see this, set $\Delta := \varepsilon^\infty(0)$ in Equation (5). A little algebra shows that the limiting spot contract rewards high output with

$$\bar{w} := \bar{\Delta} \frac{1}{1 + \beta} \frac{Q - \beta(1 - p)(p - Q)}{Q} \tag{8}$$

Under the proposed arrangement, the incentive level at all times is \bar{w} . However, since $\bar{w} < \bar{\Delta}$, the IC-constraint requires tomorrow’s incentive level after high output today to be at least $\varepsilon(\bar{w}) > \bar{w}$. Thus, the IC-constraint is always violated.

The cliff arrangement connects the myopic agency optimal contract to high-watermark contracts used in the hedge fund industry. The myopic agency model captures an important aspect of the agency problem in a hedge fund. The manager can take the long-term action of exerting effort to turn his innate skill into generating alpha. Or, the manager can take the short-term action by engaging in various “information-free activities that amount to manipulation... not based upon the production and deployment of value-relevant information about the underlying assets in the portfolio.” (Goetzmann, Ingersoll and Spiegel, 2007).

The mapping of hedge fund returns to the binary output system of my model is a little trickier. Empirical findings suggest that the returns distributions of hedge funds typically exhibit large negative skewness and high excess kurtosis. This means much of the variance is generated by negative tail events. Thus, one reasonable mapping is to let high output denote positive alpha and low output denote the negative tail event.

My paper predicts that the solution to the agency problem is the aforementioned cliff arrangement. In practice, hedge fund managers are compensated using high-watermark contracts. Essentially, a high-watermark contract pays the manager a performance-contingent fee only when cumulative performance exceeds the running maximum, that is, the high-watermark. As the manager consistently generates alpha, more funds flow in and the performance fee typically increases. This fits well with the backloaded reward schedule of the cliff arrangement.

When the negative tail event occurs, the fund’s net asset value falls well below the high-watermark and performance fees disappear. Compounding the loss, the smaller fund means that the manager takes in less from his management fee. Moreover, since the gap between current net asset value and the high-watermark is very large by definition of a negative tail event, the manager will find it difficult to close the gap quickly. This implies that when things do go south, it may be a while before the manager receives significant compensation again, much like a fall off a cliff.

It is worth noting that in particularly serious (though not all that uncommon) cases, the manager may find the gap so large that he simply gives up. That is, he closes the fund and starts anew. By doing this, the manager is effectively cushioning his fall off the cliff on the backs of his previous investors. Such a maneuver certainly dilutes the effectiveness of the high-watermark's implicit cliff arrangement. That being said, the fixed costs of starting a fund are not trivial and the manager's new fund's size and his new performance fee will likely both be smaller. Therefore, even when a manager abandons his old fund, he still experiences something comparable to a fall.⁸

3.3 Comparing with the Non-Persistent Version of the Model

The agency problems that the myopic agency model is geared toward analyzing have traditionally been dealt with in settings where actions have no persistent effects on firm performance. Thus, in an effort to assess the value of the myopic agency model and its optimal contract, a natural comparison to make is with the non-persistent version of the model.

So, suppose the true model is the myopic agency model but the principal decides to abstract away from the persistent nature of the moral hazard problem. He replaces the true model with the appropriate, simplified, non-persistent approximation.

The long-term action is relabeled as effort. Effort today is modeled to produce a probability p of generating high output today. The short-term action is relabeled shirking. Shirking today is modeled to produce a probability $1 - \beta Q$ of generating high output today. Notice, the principal replaces the short-term action's true multi-period effect on output with its present value. Shirking provides private benefit b .

Recall in the true model, I assume that the parameters are such that $\beta Q > 1 - p$. That is, the loss tomorrow associated with taking the short-term action today outweighs the present gain. In the non-persistent model, this is precisely the condition that ensures shirking generates high output with lower probability compared to effort. Therefore, Assumption (A) can still be sensibly applied and the principal can solve for the optimal contract (subject to always inducing effort). I will call this contract the *traditional contract* to distinguish it from the optimal contract of the true model.

Fortunately for the principal, sweeping the persistence of the moral hazard problem under the rug and using the traditional contract still induces the long-term action from the agent. That is, the traditional contract is incentive-compatible in the true model. Obviously, it is inefficient since it is not isomorphic to the optimal contract. The degree of inefficiency can be usefully quantified once I describe the traditional contract's structure.

⁸The case of Amaranth Advisors is instructive. Large bets placed on natural gas futures maturing in the fall and winter of 2005-06 proved catastrophic when expected price increases did not materialize. In one month, the \$9 billion dollar fund lost 70% of its value and was eventually liquidated. In the aftermath, both Nicholas Maounis, head of Amaranth, and Brian Hunter, the trader behind the large bets, attempted to start new funds. Both faced significant challenges. In 2008, Maounis started Verition Fund Management and faced fund-raising difficulties. Today, the fund is much smaller than Amaranth at its peak, and a significant portion of the fund's assets is Mr. Maounis' (Strasburg, 2010). In 2007, Brian Hunter attempted to start Solengo Capital Advisors. However, much of the initial funds raised disappeared due to potential legal action against Hunter (McLean, 2008). Eventually, the assets of Solengo were sold.

The traditional contract is stationary. The spot contract used at every date is: pay the agent $\bar{\Delta}$ if output is high and nothing if output is low. Thus the traditional contract simply repeats the optimal contract of the one shot version of the non-persistent model. The pay-to-performance sensitivity of the contract is trivially always $\bar{\Delta}$. In addition, the incentive level is also always $\bar{\Delta}$. In contrast, the optimal contract's history dependent incentive level is $\varepsilon^N(0)$ which is always below $\bar{\Delta}$, satisfying $0 = \varepsilon^0(0) < \varepsilon(0) < \varepsilon^2(0) < \dots < \varepsilon^\infty(0) = \bar{\Delta}$.

Remark. *Relative to the optimal contract, the traditional contract always over-incentivizes the agent.*

Consequently, the pay-to-performance sensitivity of the traditional contract is also too high. There is, however, a small but meaningful difference between comparing incentive levels and comparing pay-to-performance sensitivities. Notice as the agent's performance increases, the optimal contract's incentive level converges to that of the traditional contract. This convergence does not happen for pay-to-performance sensitivity. Specifically, the limiting spot contract of the optimal contract is not the spot contract of the traditional contract. Recall, Equation (5) implies that limiting reward for high output in the optimal contract is \bar{w} which is strictly smaller than $\bar{\Delta}$.

The comparison of pay-to-performance sensitivities reveals a significant difference in pay level between the optimal and traditional contracts. As the agent's performance becomes arbitrarily good, the reward he receives for producing high output in the optimal contract is still strictly smaller than the reward he would receive each time he produced high output in the traditional contract. This has important feasibility implications.

For example, suppose that the agent's private benefit b exceeds the threshold $\beta Q - (1 - p)$. Then $\bar{\Delta} > 1$ and the traditional contract pays the agent more than he produces! Clearly, there is no possible justification for inducing the long-term action from the agent. This need not be the case if the principal uses the optimal contract. Since $\bar{w} < \bar{\Delta}$, there are values of b where $\bar{\Delta} > 1$ but $\bar{w} < 1$. Moreover, even if $\bar{w} > 1$, the backloaded nature of the reward schedule implies that only a fraction of the time is the optimal contract's spot contract rewarding the agent more than he produces. One implication is that if b is not too much larger than $\beta Q - (1 - p)$, then it is still possible for the net payoff to the principal from employing the optimal contract to be larger than his payoff from giving the agent nothing and letting the agent always take the short-term action.

4 Incentive Escalation and Termination

The incentive structure is a key difference between the optimal and traditional contracts. The incentives of the optimal contract are always smaller. Also, the incentive level of the optimal contract is non-stationary, always increasing after high output. Over time, the expected incentive level $\mathbf{E}\Delta_t$ also increases and converges to $pb/(\beta Q - p(1 - p))$ as the performance indicator N settles into its stationary distribution.

Thus, the optimal contract exhibits a form of incentive escalation. It predicts that agents with histories of sustained high output should receive the highest incentives and pay-to-performance sensitivities. This escalation result complements and provides an alternative

theory of wage dynamics to Gibbons and Murphy (1992). In that paper, the authors observe that the optimal contract should emphasize total reward-to-performance sensitivity, which should factor in implicit career concerns as well as explicit pay. Escalation in explicit pay-to-performance sensitivity is then driven by the gradual disappearance of career concerns over time. Since career concerns disappear regardless of performance history, the escalation of pay-to-performance sensitivity is history independent. This contrasts with the myopic agency optimal contract that concentrates all of the escalation behind high output histories. Another difference is with the agent’s risk attitude. While the escalation result in my model is fully compatible with a risk averse agent, it is not dependent on risk aversion.

On a related note, introducing risk aversion to the agent’s utility function can produce a wealth effect that escalates a contract’s pay-to-performance sensitivity without affecting incentives. Such a dynamic should not be confused with the incentive escalation result of my paper. Incentives and pay-to-performance sensitivity are used somewhat interchangeably in the literature but they are theoretically distinct objects. Recall, incentives measure promised-value-to-performance sensitivity, which directly drives agent behavior through the IC-constraint. Fluctuations in pay-to-performance sensitivity may reflect an underlying change in incentives but may also be a symptom of other changes such as in wealth effects or career concerns.

4.1 Optimality with Termination

The incentive escalation property is the main reason I chose to not include the option to terminate in the original model. Termination is an important aspect of dynamic contracting. However, if the only agency problem is myopic agency, then giving the principal the ability to terminate the current agent and contract a new one produces a perverse optimal termination rule. When the current agent is doing well, the principal wants to terminate him due to the incentive escalation that occurs after high output and recontract with a new agent at a lower incentive level. Conversely, when the current agent is doing poorly, the principal paradoxically wants to keep him. This unwillingness to terminate the agent when he is doing poorly is a consequence of a severance pay result that I will discuss shortly. Ultimately, it boils down, again, to the incentive escalation property.

Thus, even though these termination tendencies are perverse, they are rational responses given the nature of this particular second-best setting. Of course, by focusing on myopic agency, the model ignores important facets of the principal-agent relationship such as adverse selection and power dynamics. These considerations will cause the principal to strongly reconsider how to optimally terminate the agent. Instead of adding these other elements to the model and letting the termination choice arise endogenously, I will, for the sake of simplicity, fix a “natural” benchmark termination rule. This rule will essentially serve as a proxy for those unmodelled elements. I will call the resulting optimal contract under the exogenous termination rule, *the optimal contract with termination*.

The rule is: terminate the agent if and only if low output is produced. This rule is the simplest one that subscribes to the more intuitive and empirically justified idea that the agent should be fired for poor not good performance. Such a rule could be optimal if, for example, a low output indicates there is a high likelihood that the agent is of low type or

has become incompatible with the firm.

It's also worth noting, the rule is the endogenously optimal termination rule of the non-persistent version of the myopic agency model. Imposing this rule in the myopic agency model maintains comparability between the optimal and traditional contracts when the option to terminate is added. This feature is useful as it will allow me to demonstrate that all of the advantages the optimal contract has over the traditional alternative in the no termination setting are more or less preserved in the termination setting.

The derivation of the optimal contract with termination mirrors the original analysis. The state variable is still Δ , and the Bellman equation characterizing optimality is only slightly changed: $V(\Delta) = \max\{\beta V(\varepsilon(\Delta)) - (1-p)\Delta, p\Delta\}$. The $\beta V(0)$ is deleted from the second argument because for the incumbent agent, the continuation value following low output/termination is his outside option, which is zero. Note, I am implicitly assuming that the contract's payments to the agent end upon termination. In principle, since actions are persistent, disallowing post-termination compensation is with loss of generality. However, this is not the case under the chosen termination rule. Recall the IC-constraint does not involve any future incentives following low output today. Thus, any type of post-termination performance sensitive pay, which by definition of the termination rule can only occur after a low output date, would serve no incentive purpose. Therefore, all such pay can be aggregated into a lump sum payment to the agent upon termination.

Mirroring the previous approach, the new Bellman equation can be solved and the optimal contract with termination can be explicitly characterized. I now describe a representative $V(\Delta)$ and analyze the corresponding optimal contract with termination.

Suppose the parameters of the model satisfy $\beta \leq \frac{Q-(1-p)p}{pQ}$. Then $V(\Delta)$ is a two-piece piecewise linear function. The two slopes are $m(p)$ and p just like in the no termination representative example. $V(0) = pb/Q$ and $x_1 = pb/(Q - p(1-p))$. Let w_t denote the agent's salary in year t if high output is produced and he is retained. Let S_t denote the agent's severance pay in year t if low output is produced and he is terminated. Then the optimal contract's salary, severance pay and incentive level are as follows:

- Salary: $w_0 = 0$. For $t > 0$,

$$0 < w_1 < \dots < w_t = \varepsilon^t(0) \left[\frac{Q - (1-p)p}{Q} \right] - \frac{pb}{Q} < \dots < w_\infty = \bar{\Delta}(1-p\beta)$$

- Severance pay: $S_0 = \frac{pb}{Q}$. For $t > 0$, $S_t = 0$.
- Incentive level: $0 = \Delta_0 < \dots < \Delta_t = \varepsilon^t(0) < \dots < \Delta_\infty = \bar{\Delta}$.

4.2 Upward Sloping Pay Scale

When the agent is junior, his salary is zero.⁹ After the junior period, salary becomes nonzero and increases with tenure. Eventually, as the agent becomes senior, his salary approaches

⁹In this example, junior means $t = 0$. In general, the junior period may last longer but not less than one date.

the steady state w_∞ . Thus, the agent faces a somewhat typical looking upward sloping pay scale. In contrast, the traditional contract with termination simply pays the fixed salary w_∞ .

Remark. *The agent's salary is increasing over time in the optimal contract with termination. However, regardless of his seniority, the salary he receives under the optimal contract with termination is always strictly smaller than the fixed salary he would receive in the traditional contract with termination.*

In general, the salary comparison between the optimal and traditional contracts with termination approximate the previous pay comparison between the no termination contracts. I will not say anything more about this.

The upward sloping pay scale w_t faced by the agent in the optimal contract with termination is a direct result of the incentive escalation property. Upward sloping pay scales are important features of many real life employment contracts. Unlike the approach taken in this paper, the traditional incentives-based arguments for upward sloping pay scales begin with the following backloading idea: In general, it is not optimal to tie pay, on a date-by-date basis, to marginal product or the flow payoff of the outside option. By delaying a payment, the principal can efficiently reuse the threat of withholding that payment to motivate the agent across multiple dates. Delaying payment also keeps the agent's promised value away from his participation constraint. This is useful in settings where the participation constraint can be binding, thereby forcing the principal to trigger an inefficient early termination. Notice this is not a relevant concern for my myopic agency model. What the idea suggests is that the principal should pay the agent only on the last date. Technically speaking, loading all the payments onto the last date does produce an upward sloping pay scale, albeit an extreme one. To get a more realistic shape, a counterbalance is typically introduced.

For example, Lazear (1981) argues that the principal also has incentive problems. In particular, there is a risk that he may renege on payments promised to the agent. In this situation, loading everything onto the last date is far too risky and so the agent's pay increases more gradually over time. Another argument used by many authors appeals to risk aversion. With a sufficiently risk averse agent, it is simply far too expensive to exclusively use the last date's pay to generate the necessary incentives that motivate the agent throughout the entire contract. Again, a more smoothly upward sloping pay scale results.

In my model, the principal is fully committed and the agent is risk neutral. Therefore, myopic agency and incentive escalation combine to provide a fully independent reason for having an upward sloping pay scale.

4.3 Severance Pay

The incentive escalation property provides a simple intuition for severance pay - potentially even large, mostly lump sum, severance pay. Suppose after some low output event it is optimal to terminate the agent. Facing such a situation, the principal may still want to pay the agent a nontrivial amount simply to dull incentives today. Why might the principal want to dull incentives today? By dulling incentives the principal need not escalate incentives as

much tomorrow following high output. This is a good thing, because incentive escalation, which is a necessary tool used to combat the agency problem, is expensive.

Therefore, severance pay for poor performance is a mechanism to temper today's incentives so as to preserve the future viability of the principal-agent relationship should it continue.

The nontrivial S_0 of the optimal contract with termination demonstrates that the severance pay intuition works in practice. If the principal deletes S_0 then one of two things will happen: either the agent takes the short-term action or the principal must increase the incentives at all dates $t > 0$ to compensate for the increased date 0 incentives. Neither is palatable.

The severance pay S_0 is proportional to the private benefit b . Therefore, its value relative to the surplus generated by the firm is ambiguous. What is not ambiguous is its value relative to the contract. Since the contract's initial incentive level is zero, the severance pay S_0 is equal in value to all of the agent's future expected earnings had he generated high output today and avoided termination. As this example demonstrates,

Remark. *Even relatively massive, lump sum severance pay can be supported as optimal arrangements.*

That being said, the example should not be interpreted as a blanket justification of golden parachutes. The maximum possible size of the severance pay today is proportional to the degree to which today's incentive level can be compressed. In the optimal contract with termination, there is no bound on incentive compression at date 0. But if in addition, there was, say, a non-persistent moral hazard problem or some other constraint on the degree of incentive compression, the maximal severance pay would be smaller. Moreover, the maximal severance pay need not always be the optimal severance pay. I chose to highlight an example of an optimal severance pay that was both maximal and purely lump sum mostly because this type of severance pay has traditionally been the hardest to justify as efficient. In general, the myopic agency intuition is compatible with severance packages both large and small, immediate and vested, lump sum and performance-contingent.

One curious feature of the optimal contract with termination is that severance pay is decreasing with tenure. In fact, if tenure is sufficiently long, the optimal contract with termination gives no severance pay.¹⁰ This unusual dynamic is an artifact of a simplifying assumption of the model: the agent has a worthless outside option. If, instead, the agent's outside option increases with tenure, which seems likely if the labor pool is heterogeneous with respect to ability, then it's possible for severance pay to be also increasing.

The intuition is more subtle than it may appear, so let me be specific. If the outside option is allowed to be nonzero, then it becomes a potentially binding constraint in the model. In general, a binding participation constraint will induce payment to the agent even upon termination. For a trivial demonstration, suppose the agent has log utility. Then at each date and after each history, the agent must be paid something. In particular, he must be paid even when he's terminated. This participation-constraint induced "severance pay

¹⁰In this example, sufficiently long means $t > 0$. In general, the initial period with severance pay may last longer but not less than one date.

intuition” is *not* what I am referring to. By severance pay, I implicitly mean a payment upon termination that is given even though it is not necessary in the sense of satisfying the agent’s ex-ante or interim participation constraint. The increasing outside options intuition to which I am referring can be described as follows: At date t , suppose high output leads to an increase in payment at date $t + 1$ driven by the assumed increase in the outside option. Then it may be optimal to have severance pay at time t , even if giving nothing at time t doesn’t break the time t participation constraint. The reasoning is the same as before, reflecting the potential benefits of dulling incentives. Clearly, if the outside option increases rapidly enough, the arrangement will pull severance pay upwards as well.

In contrast, the traditional contract with termination exhibits no severance pay. The downside of increasing today’s incentives that drives the severance pay intuition in the myopic agency setting is simply not applicable in settings with non-persistent moral hazard. The IC constraint in non-persistent settings is a lower bound on each date’s incentives. Deleting severance pay for low output only further ensures that the IC constraint is not violated. This intuition explains why the dynamic contracting literature traditionally has had difficulty explaining severance pay.

The severance pay result in this paper is not driven simply by the fact that actions have persistent effects. It is worth distinguishing between the intuition described in this paper coming from myopic agency in particular and the alternate motivation for severance pay coming from action persistence in general. The latter idea simply argues that since there is persistence, actions taken by the agent continue to have effects on the firm post-termination. Therefore, the principal may want to have part of the agent’s compensation postponed to after termination - hence severance pay. This is clearly not the same idea as the one relying on myopic agency and incentive escalation. The intuition for severance pay in this paper, for example, can also explain lump sum severance pay, which is something the action persistence intuition cannot explain.

5 Conclusion

Short-termism is a major component of many managerial agency problems. This paper investigates optimal contracting when a manager can take hidden short-term actions that hurt the future health of the firm. Like in many real-life settings, the short-term action in this model boosts performance today. This temporarily masks the inferiority of the short-term action and creates a tricky contracting setting where simply rewarding high output is no longer guaranteed to eliminate the agency problem. In this setting, I show that the optimal contract always inducing the long-term action differs in significant ways from traditional dynamic optimal contracts. The myopic agency optimal contract values sustained high output. Its incentive level is non-stationary, escalating over time and after high output while remaining strictly below the incentive level of the traditional contract. Also, severance pay may accompany termination for low output. In addition, the myopic agency optimal contract provides new perspectives on high-watermark contracts and upward sloping pay scales. More generally, the paper establishes a framework that can be used to model a variety of agency problems where the assumption that actions have no persistent effects is flawed.

6 Appendix

Proof of Lemma 1. Fix a contract that calls for the agent to always take the long-term action. Consider the diagram below representing today's pay and tomorrow's "ex-post" promised values following a history h leading up through yesterday. Promised values are calculated with respect to the measure generated by always taking the long-term action.

$$\begin{array}{rcl}
 & & w(h11) + \delta W(h11) \\
 w(h1) & & \\
 & & w(h10) + \delta W(h10) \\
 \\
 & & w(h01) + \delta W(h01) \\
 w(h0) & & \\
 & & w(h00) + \delta W(h00)
 \end{array}$$

Suppose the agent decides to commit a one-shot deviation to the short-term action today. Then his payoff is $b + w(h1) + \beta(\pi(p(w(h11) + \delta W(h11)) + (1 - p)(w(h10) + \delta W(h10))) + (1 - \pi)(w(h10) + \delta W(h10)))$. Letting $W(h1)$ and $W(h0)$ denote the agent's promised values tomorrow following high and low output today (again, calculated under the measure generated by always taking the long-term action), the utility from deviation can be rewritten as $b + (w(h1) + \delta W(h1)) - \beta((1 - \pi)p((w(h11) + \delta W(h11)) - (w(h10) + \delta W(h10)))) = b + (w(h1) + \delta W(h1)) - \beta Q \Delta(h1)$. Incentive compatibility requires that

$$p(w(h1) + \delta W(h1)) + (1 - p)(w(h0) + \delta W(h0)) \geq b + (w(h1) + \delta W(h1)) - \beta Q \Delta(h1)$$

which, upon rearrangement, is equivalent to

$$\Delta(h1) \geq \frac{(1 - p)\Delta(h) + b}{\beta Q}$$

Thus the proposed IC-constraint is a necessary condition ensuring that one-shot deviations from always taking the long-term action are suboptimal. I now show sufficiency by proving that if one-shot deviations are suboptimal then all deviations are suboptimal.

First note, if after some history the agent is better off employing a deviation strategy, then he is better employing a deviation strategy that only involves deviating in a finite number of dates. This is due to discounting. This observation allows me to prove sufficiency using induction.

So fix a contract that always calls for the long-term action and satisfies the proposed IC-constraint. Suppose there are no profitable T -length deviations. Now, suppose on the contrary, there exists a history h such that following h there exists a profitable $T + 1$ -length deviation. If this deviation does not involve deviating right away, then it is in fact, a T -length deviation. Contradiction. So suppose the deviation does involve deviating right away. Then the agent's payoff following h is $b + w(h1) + \beta(\pi U(D(h1)) + (1 - \pi)(w(h10) + \beta U(D(h10))))$ where $U(D(h1))$, $U(D(h10))$ are the payoffs from employing the continuations

$D(h1)$, $D(h10)$ of the deviation strategy after histories $h1$ and $h10$. By induction, this payoff is weakly less than $b + w(h1) + \beta(\pi(p(w(h11) + \delta W(h11)) + (1 - p)(w(h10) + \delta W(h10))) + (1 - \pi)(w(h10) + \beta(p(w(h101) + \delta W(h101)) + (1 - p)(w(h100) + \delta W(h100)))) = b + w(h1) + \beta(\pi(p(w(h11) + \delta W(h11)) + (1 - p)(w(h10) + \delta W(h10))) + (1 - \pi)(w(h10) + \delta W(h10))) \leq p(w(h1) + \delta W(h1)) + (1 - p)(w(h0) + \delta W(h0))$.

The key step in the proof is to realize that the payoffs from employing the continuations of the deviation strategy are the same regardless of the initial deviation at history h . This is what allows the inductive step to go through. \square

Proof of Lemma 2 and Theorem 1.

Step 1. $V(\Delta)$ is convex and the optimal contract is randomization-proof.

Fix $\lambda \in (0, 1)$ and Δ_1 - and Δ_2 -contracts with costs $C(\Delta_1)$ and $C(\Delta_2)$ and payments $w_{\Delta_1}(h)$ and $w_{\Delta_2}(h)$ after each history h . Then the contract that pays $\lambda w_{\Delta_1}(h) + (1 - \lambda)w_{\Delta_2}(h)$ is a $\lambda\Delta_1 + (1 - \lambda)\Delta_2$ -contract with cost $\lambda C(\Delta_1) + (1 - \lambda)C(\Delta_2)$.

To show that the optimal contract is randomization-proof, consider a contract with possibly random payments to the agent based on a sequence of random signals $\theta_0, \theta_1 \dots$ where the payments at date t can depend on $\{\theta_0, \dots, \theta_t\}$ in addition to h_t . Assume the action at date t is chosen before the realization of θ_t . Define $W_t(h_{t-1}, \theta_0, \dots, \theta_{t-1})$ as the appropriate generalization of $W_t(h_{t-1})$ where the expectation is also taken over all possible present and future signal realizations $\theta_t, \theta_{t+1} \dots$ in addition to all possible future histories. We can then define $\Delta_t(h_{t-1}, \theta_0, \dots, \theta_{t-1})$ as the appropriate generalization of $\Delta_t(h_{t-1})$. A slight generalization of the proof of Lemma 1 shows that the IC-constraint with public randomization is:

$$\mathbf{E}_{\theta_t} \Delta(h_{t-1}, \theta_0, \dots, \theta_t) \geq \frac{(1 - p)\Delta(h_{t-1}, \theta_0, \dots, \theta_{t-1}) + b}{\beta Q}$$

If the optimal contract were not randomization-proof, then there would exist some contract with randomization only in a finite number of dates that was less costly. Thus, it suffices to show for every n , given any contract with randomization only in the first n dates and initial incentive level Δ , there exists a non-random Δ -contract with equal cost. Here, “non-random” only means it is not a function of the public randomization variables.

The proof is by induction. The $n = 1$ case: Consider a contract that only randomizes over θ_0 and has initial incentive level Δ . Let $w_0(h_0, \theta_0)$ be the random date 0 payments of the contract. Define $w_0(h_0) := \mathbf{E}_{\theta_0} w_0(h_0, \theta_0)$. Given a realization (h_0, θ_0) , the date 1 continuation contract is by assumption some non-random $\Delta(h_0, \theta_0)$ -contract with some cost $C(\Delta(h_0, \theta_0))$. Define $\Delta(h_0) := \mathbf{E}_{\theta_0} \Delta(h_0, \theta_0)$ and $C(\Delta(h_0)) := \mathbf{E}_{\theta_0} C(\Delta(h_0, \theta_0))$. The proof of the convexity of V shows that there exists a non-random $\Delta(h_0)$ -contract with cost $C(\Delta(h_0))$.

Consider the following non-random contract: pay $w_0(h_0)$ at date 0 after history h_0 , followed by the $\Delta(h_0)$ -continuation contract. By construction, it has the same cost as the original contract with randomization and is incentive-compatible for all dates $t \geq 1$. All that is left to show is incentive-compatibility at date 0:

$$\Delta(1) = \mathbf{E}_{\theta_0} \Delta(1, \theta_0) \geq \frac{(1 - p)\Delta + b}{\beta Q}$$

So the $n = 1$ case is proved.

Now suppose it is proved for $n = N$. Consider a contract with randomization only in the first $N+1$ dates. For every realization (h_0, θ_0) , consider the date 1 continuation contract with randomization in at most N dates. By assumption, it can be replaced with a non-random $\Delta(h_0, \theta_0)$ -contract. And now we are back to the $n = 1$ case.

Step 2. $\Delta^* < \varepsilon(0)$ where $\Delta^* \in \arg \min V(\Delta)$.

Suppose not. Then consider the following contract: at date 0, pay the agent nothing; at date 1 give the agent the optimal Δ^* -contract. By assumption, this contract is incentive-compatible. It is a 0-contract by construction. It has cost $\beta V(\Delta^*) < V(\Delta^*) = \min V(\Delta)$. Contradiction.

Step 3. $V(\Delta) = \max\{\beta V(\varepsilon(\Delta)) - (1-p)\Delta, p\Delta + \beta V(\Delta^*)\}$.

Consider the following sequence of panels:

$$\begin{array}{ccc}
 w(1) & C(\Delta(1)) & w(1) + \beta(C(\Delta(1)) - V(\varepsilon(\Delta))) & V(\varepsilon(\Delta)) \\
 & & \longrightarrow & \longrightarrow \\
 w(0) & C(\Delta(0)) & w(0) + \beta(C(\Delta(0)) - V(\Delta^*)) & V(\Delta^*)
 \end{array}$$

The first panel depicts a generic Δ -contract: it's date 0 high and low output payments as well it's date 1 continuation contracts' costs. The second panel depicts the first step of a two step cost-decreasing transformation of the Δ -contract. The continuation contracts are changed to be the least costly ones subject to the date 0 incentive constraint and the resulting distortion in promised value is balanced out by increasing the date 0 payments by appropriate amounts. That the date 1 continuation contract following high output is the optimal $\varepsilon(\Delta)$ -contract follows from Step 1 and Step 2. Notice, the resulting contract in the second panel is also a Δ -contract; it has the same cost as the first panel's Δ -contract; and it has weakly larger date 0 payments.

$$\begin{array}{ccc}
 \left(\Delta - \beta[V(\varepsilon(\Delta)) - V(\Delta^*)]\right)^+ & V(\varepsilon(\Delta)) & \\
 & \implies & \\
 \left(\beta[V(\varepsilon(\Delta)) - V(\Delta^*)] - \Delta\right)^+ & V(\Delta^*) &
 \end{array}$$

The third panel depicts the second step of the two step cost-decreasing transformation. The common portion of the date 0 high and low output payments from the previous panel's Δ -contract has been factored out. This transformation weakly decreases the cost of the contract while maintaining its property of being a Δ -contract. The date 0 high and low output payments are now the smallest possible ones for a Δ -contract. Thus, the contract represented by the third panel is in fact the optimal Δ -contract which, by definition, has cost $V(\Delta)$. This cost can be explicitly computed using the diagram, resulting in the following piecewise equation for $V(\Delta)$:

$$V(\Delta) = \begin{cases} \beta V(\varepsilon(\Delta)) - (1-p)\Delta & \text{if } \left(\beta[V(\varepsilon(\Delta)) - V(\Delta^*)] - \Delta \right)^+ \geq 0 \\ p\Delta + \beta V(\Delta^*) & \text{if } \left(\Delta - \beta[V(\varepsilon(\Delta)) - V(\Delta^*)] \right)^+ \geq 0 \end{cases}$$

A little algebra then shows that this piecewise equation can be recast as the following Bellman equation:

$$V(\Delta) = \max\{\beta V(\varepsilon(\Delta)) - (1-p)\Delta, p\Delta + \beta V(\Delta^*)\}$$

Note this Bellman equation differs slightly from Equation (3). For the rest of the proof, I will call this Bellman equation the *weak Bellman equation* and Equation (3) the *strong Bellman equation*. The weak Bellman equation becomes the strong Bellman equation if $V(\Delta)$ is weakly increasing.

Step 4. $V(\Delta)$ is a weakly increasing, piecewise linear function.

Let \mathcal{C} be the space of all weakly convex functions $f(x)$ defined on $[0, b/(\beta Q - (1-p))]$ satisfying $0 \leq f(x) \leq pb/[(1-\beta)Q] + px$ and $-Q(1-p)/(Q - (1-p)) \leq f'(x) \leq f'(x) \leq p$ for all x . Under the L^∞ norm, \mathcal{C} is a compact Banach space.

Define the operator $Tf(x) := \max\{\beta f(\varepsilon(x)) - (1-p)x, px + \beta \min\{f(x)\}\}$ coming from the weak Bellman equation.

Step 4a. T is a contraction of \mathcal{C} .

We first need to show that T is an endomorphism of \mathcal{C} . Define operators T_a and T_b where $T_a f(x) := \beta f(\varepsilon(x)) - (1-p)x$ and $T_b f(x) := px + \beta \min\{f(x)\}$. If f is weakly convex, then clearly so are $T_a f$, $T_b f$. Then $Tf = \max\{T_a f, T_b f\}$ is weakly convex. If f is nonnegative, then clearly so is $T_b f$ and therefore Tf . If $f'(x) \geq -Q(1-p)/(Q - (1-p))$ for all x , then $T_a f'(x) = \beta f'(\varepsilon(x)) - (1-p) \geq \beta(-Q(1-p)/(Q - (1-p)))(1-p)/(\beta Q - (1-p)) = -Q(1-p)/(Q - (1-p))$ for all x . Also $T_b f'(x) = p$. Therefore $Tf'(x) \geq -Q(1-p)/(Q - (1-p))$ for all x . A similar argument shows that if $f'(x) \leq p$ for all x then $Tf'(x) \leq p$ for all x . Lastly, suppose $f(x) \leq pb/[(1-\beta)Q] + px$ for all x . Then $T_a f(x) = \beta f(\varepsilon(x)) - (1-p)x \leq \beta pb/[(1-\beta)Q] + \beta p[(1-p)x + b]/(\beta Q - (1-p)) = (\beta pb)/Q \cdot [1/(1-\beta) + 1/\beta] + [p(1-p)/Q - (1-p)]x \leq pb/[(1-\beta)Q] + px$. Also $T_b f(x) < pb/[(1-\beta)Q] + px$ for all x . Thus $Tf(x) \leq pb/[(1-\beta)Q] + px$ for all x . I have proved $T(\mathcal{C}) \subset \mathcal{C}$.

Let $f, g \in \mathcal{C}$. For any x , $|Tf(x) - Tg(x)| \leq \max\{|T_a f(x) - T_a g(x)|, |T_b f(x) - T_b g(x)|\}$. Since both T_a and T_b are contractions on \mathcal{C} , T is a contraction on \mathcal{C} . Therefore, a unique fixed point of \mathcal{C} under T is assured.

Step 4b. $V(\Delta)$ is weakly increasing and therefore, satisfies the strong Bellman equation.

In this sub-step, I will affix the superscript Q to \mathcal{C} and T to emphasize that these objects change as I change the model parameter Q .

Consider a realization of the model with $Q = p$, space \mathcal{C}^p and operator T^p . It is easy to show that for all n , $(T^p)^n 0$ is a convex two-piece piecewise linear function where the first piece is flat and the second piece has slope p . In particular, $(T^p)^n 0$ is nondecreasing for all n and therefore, so is the unique fixed point of T^p .

Now consider any model realization with $Q < p$ with corresponding \mathcal{C}^Q and T^Q . Let $f^Q \in \mathcal{C}^Q$ and $f^p \in \mathcal{C}^p$. Then I say $f^Q > f^p$ if and only if on the smaller domain of \mathcal{C}^p functions, $f^Q - f^p$ is nondecreasing, nonnegative and not identically zero.

It is easy to show that if f^Q and f^p are both nondecreasing and $f^Q \geq f^p$, then $T^Q f^Q > T^p f^p$. Thus $(T^Q)^n 0$ is strictly increasing for all $n > 1$ and therefore, so is the unique fixed point of T^Q .

Step 4c. $V(\Delta)$ is a piecewise linear function.

Fix a $\Delta \in (0, b/(\beta Q - (1 - p)))$. Define $N \geq 1$ to be the unique integer satisfying $\varepsilon^{-N}(\Delta) \leq 0 < \varepsilon^{-N+1}(\Delta)$. Define $m(x) := (1 - p)x/Q - (1 - p)$. Construct the following piecewise linear function v piece by piece starting from the left:

$$v(x) = \begin{cases} v(0) + m^N(p)x & x \in (0, \varepsilon^{-N+1}(\Delta)] \\ v(\varepsilon^{-N+1}(\Delta)) + m^{N-1}(p)(x - v(\varepsilon^{-N+1}(\Delta))) & x \in (\varepsilon^{-N+1}(\Delta), \varepsilon^{-N+2}(\Delta)] \\ \dots & \\ v(\varepsilon^{-1}(\Delta)) + m(p)(x - v(\varepsilon^{-1}(\Delta))) & x \in (\varepsilon^{-1}(\Delta), \Delta] \\ v(\Delta) + p(x - v(\Delta)) & x > \Delta \end{cases}$$

where $v(0)$ is (uniquely) defined to satisfy $v(\Delta) = p\Delta + \beta \min\{v\}$. Notice, it must be that $v(0) > 0$.

As a function of Δ , $d(\Delta) := v(0) - v(\varepsilon(0))$ is strictly increasing. For Δ sufficiently close to 0, $d(\Delta) < 0$. For Δ sufficiently large - it suffices for Δ to be large enough so that $m^{N-1}(p) \leq 0$, $d(\Delta) > 0$. Let Δ^* satisfy $d(\Delta^*) = 0$. Then the corresponding function v^* , when restricted to \mathcal{C} , is the unique fixed point of T . So $V \equiv v^*$. \square

References

- [1] Abreu, D., D. Pearce and E. Stacchetti (1990) "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* Vol. 58, pp. 1041-1063
- [2] Edmans, A., X. Gabaix, T. Sadzik and Y. Sannikov (2012) "Dynamic CEO Compensation," *The Journal of Finance* Vol. 67, pp. 1603-1647
- [3] Fernandes, A. and C. Phelan (2000) "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory* Vol. 91, pp. 223-247
- [4] Gibbons, R. and K. J. Murphy (1992) "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy* Vol. 100, pp. 468-505
- [5] Goetzmann, W., J. Ingersoll and M. Spiegel (2007) "Portfolio Performance Manipulation and Manipulation-proof Performance Measures," *Review of Financial Studies* Vol. 20, pp. 1503-1546

- [6] Green, E. (1987) "Lending and the Smoothing of Uninsurable Income," in *Contractual Arrangements for Intertemporal Trade*, ed. by E.C. Prescott and N. Wallace. Minneapolis: University of Minnesota Press pp. 3 - 25
- [7] Holmstrom, B. and P. Milgrom (1987) "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica* Vol. 55, pp. 303 - 328
- [8] Holmstrom, B. and P. Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, & Organization* Vol. 7, pp. 24-52
- [9] Jarque, A. (2010) "Repeated Moral Hazard with Effort Persistence," *Journal of Economic Theory* Vol. 145, pp. 2412-2423
- [10] Lazear, E. (1981) "Agency, Earning Profiles, Productivity, and Hours Restrictions," *The American Economic Review* Vol. 71, pp. 606 - 620
- [11] Manso, G. (2012) "Motivating Innovation," forthcoming in *The Journal of Finance*
- [12] McLean, B. "The Man Who Lost \$6 Billion," *Fortune*, July 8, 2008
- [13] Radner, R. (1981) "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica* Vol. 49, pp. 1127-1148
- [14] Radner, R. (1985) "Repeated Principal-Agent Games with Discounting," *Econometrica* Vol. 53, pp. 1173-1198
- [15] Spear, S.E. and S. Srivastava (1987) "On Repeated Moral Hazard with Discounting," *Review of Economic Studies* Vol. 54, pp. 599 - 617
- [16] Stein, J. (1988) "Takeover Threats and Managerial Myopia," *Journal of Political Economy* Vol. 96, pp. 61 - 80
- [17] Stein, J. (1989) "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics* Vol. 104, pp. 655 - 69
- [18] Strasburg, J. "Former Amaranth Chief Returns to the Stage," *The Wall Street Journal*, June 26th, 2010
- [19] Varas, F. (2013) "Contracting Timely Delivery with Hard to Verify Quality," working paper